



Variable Radii Poisson-Disk Sampling

**Scott A. Mitchell, Alexander Rand,
Mohamed S. Ebeida, Chandrajit Bajaj**

www.cs.sandia.gov/~samitch

(or Google Mitchell Sandia)


24th Canadian Conference on Computational Geometry

8-10 August 2012

Session 4B

Thursday Aug 9

11:30-11:50





Goal – convince you

- **There is still interesting Computational Geometry work for generating separate-yet-dense point sets**
 - Delaunay Refinement (DR) doesn't solve everything
- **Poisson-disk output has some advantages**
 - Graphics cares
 - Fracture mechanics cares
- **Even though slower than deterministic DR**
- **Slightly different than sphere packings**



Outline

- **Maximal Poisson-Disk Sampling (MPS) – what is it, why do we care**
 - Graphics apps
 - Simulations
- **Our prior results for MPS points, Voronoi and Delaunay meshes**
 - Sites may encroach on boundary, not dual of a body-fitted tetrahedralization
- **Spatially varying radii**
 - Lipschitz conditions
- **Motivation: MPS spectrum vs. blue noise**
- **Two-radii MPS definition**
 - Random refinements
- **MPS output vs. Delaunay Refinement (DR)**
 - PSA Spectrum
 - Angle spectra vs. DR, Edge length vs. DR



Maximal Poisson-Disk Sampling


- What is MPS?
 - Dart-throwing
 - Insert random points into a domain, build set X
 - With the “Poisson” process

Empty disk: $\forall x_i, x_j \in X, x_i \neq x_j : \|x_i - x_j\| \geq r$

Bias-free: $\forall x_i \in X, \forall \Omega \subset \mathcal{D}_{i-1} :$

$$P(x_i \in \Omega) = \frac{\text{Area}(\Omega)}{\text{Area}(\mathcal{D}_{i-1})}$$

Maximal: $\forall x \in \mathcal{D}, \exists x_i \in X : \|x - x_i\| < r$






MPS a.k.a.

- Statistical processes
 - Hard-core Strauss disc processes
 - Non-overlap: inhibition distance r_1
 - cover domain: disc radius r_2
- Geostatistics focus is inverse problem
 - Given satellite pictures (non-maximal dist.)
 - How many trees are there?
 - How much lumber?
 - Trees in a forest
 - Points are trunks



New Mexico “Forest”

$r_1 > r_2$




Disks are canopy or
separation distance




British Columbia Forest

$r_1 < r_2$



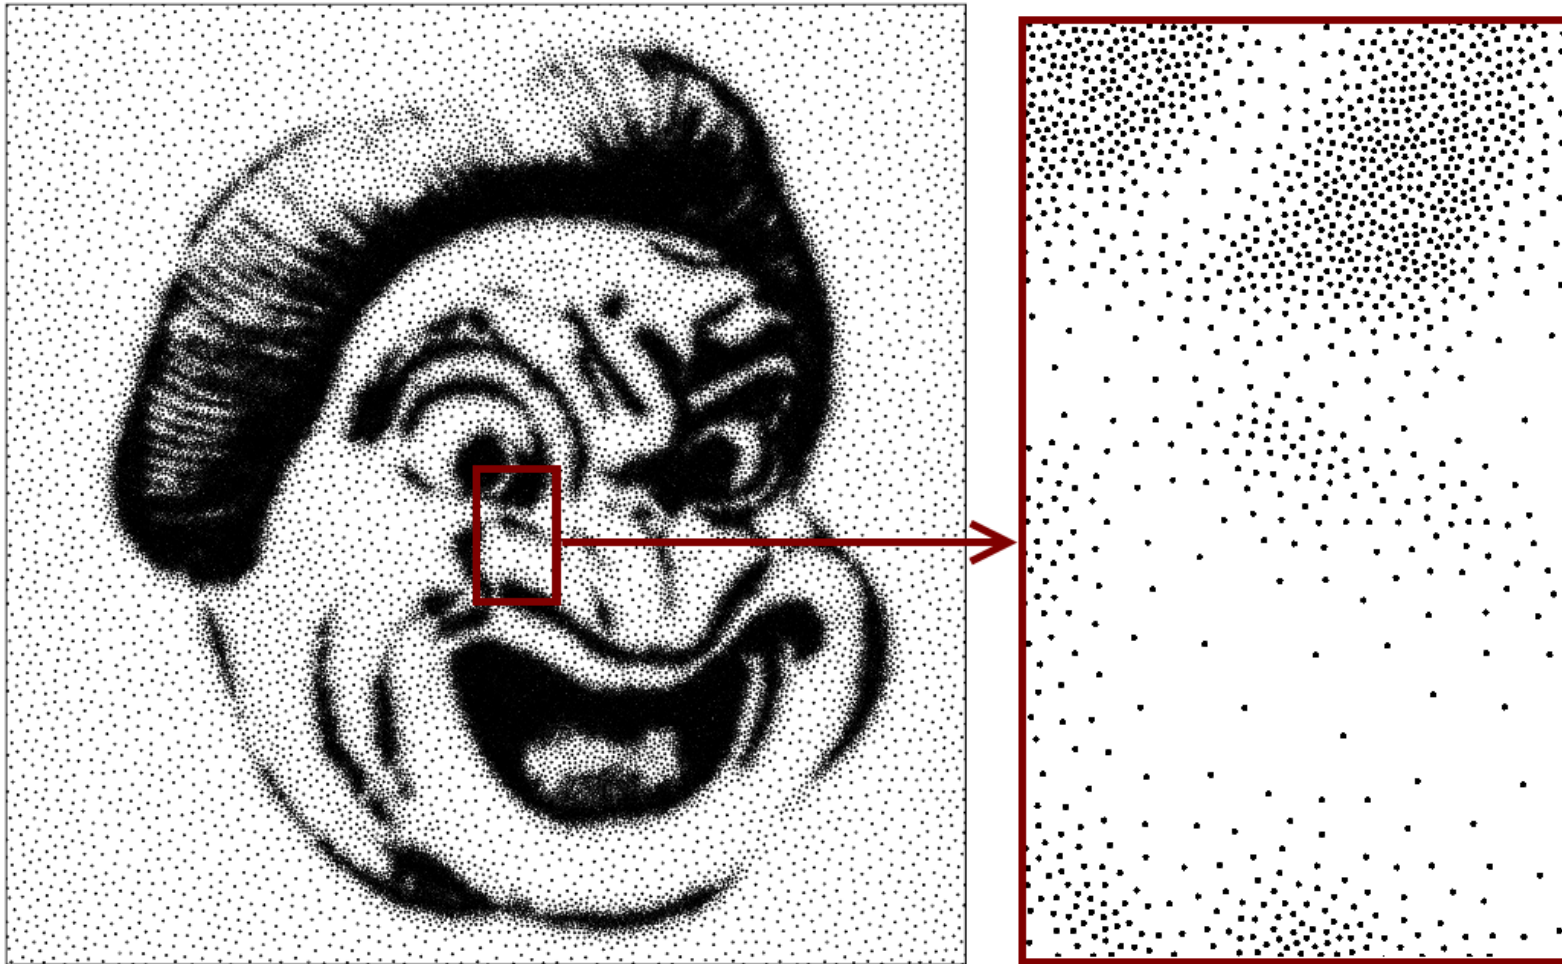
- Random sphere packing
 - $r/2$ -disks non-overlap
 - MPS = random sequential adsorption
 - (atoms in a liquid, crystal)





Motivation from old-school Static Graphics

- **Stippling:** images from dots, as newsprint
(results from this paper)



Motivating from Modern Graphics: (Brush) Stroke-Based Rendering

- CG artistic effect to mimic physical media
- Images from Aaron Hertzmann, Stroke-Based Rendering



Source photo




Painted version




Final rendering

Definition: A **stroke** is a data structure that can be rendered in the image plane. A **stroke model** is a parametric description of strokes, so that different parameter settings produce different stroke positions and appearances.


For example, one form of stippling uses a very simple stroke model:



Stippling stroke model Individual strokes (stipples)



Vector field




Final rendering




Motivating from Modern Graphics: Texture Synthesis

- Real-time environment exploration. **Games! Movies!**
- Algorithm to create output image from input sample
 - Arbitrary size
 - Similar to input
 - No visible seams, blocks
 - No visible, regular repeated patterns



examples from wikipedia:



Spaghetti
Li Yi Wei
SIGGRAPH 2011



What is MPS good for?

Sandia cares about Games and Movies? training...


- **Physics simulations – why SNL paid for year 1-2 ☺**

- Voronoi mesh, cell = points closest to a sample
- Fractures occur on Voronoi cell boundaries
 - Mesh variation \subset material strength variation
 - CVT, regular lattices give unrealistic cracks



- **Unbiased sampling gives realistic cracks**

- **Ensembles of simulations**

- **Domains: non-convex, internal boundaries**




Seismic Simulations
maximal helps Δ quality



Fracture Simulations

Courtesy of
Joe Bishop (SNL)










Prior Results

- Many fast Graphics algorithms that modified the process slightly, or the termination criteria
- First $E(n \log n)$ algorithm with provably correct output
 - Efficient Maximal Poisson-Disk Sampling, Ebeida, Patney, Mitchell, Davidson, Knupp, Owens, SIGGRAPH 2011
- Simpler, less memory, provably correct, faster in practice but no run-time proof
 - A Simple Algorithm for Maximal Poisson-Disk Sampling in High Dimensions, Ebeida, Mitchell, Patney, Davidson, Owens Eurographics 2012
- Voronoi Meshes
 - Sites interior, close to domain boundary are OK, not the dual of a body-fitted Delaunay Mesh
 - Uniform Random Voronoi Meshes Ebeida, Mitchell IMR 2011
- Delaunay Meshes
 - Protect boundary with random balls
 - Efficient and Good Delaunay Meshes from Random Points Ebeida, Mitchell, Davidson, Patney, Knupp, Owens SIAM GD/SPM 2011

UC DAVIS
UNIVERSITY OF CALIFORNIA







First Contribution


- Uniform, static literature pretty rigorous
 - Graphics papers with heuristics for sampling curved surfaces, non-uniformly
 - unknown or unstated Lipschitz criteria, neighbor datastructures that sometimes blow-up in practice
- Lipschitz conditions for spatially varying radii function
 - Reasons and proof techniques as in Delaunay Refinement



How fast can radii vary?

- If varies slowly
 - bounded # neighbors for disk conflict checks \leftrightarrow bounded-angle DT
- If shrink too fast
 - Unbounded # neighbors
 - Infinite run-time
 - Zero angles in triangulation






Q. How fast can it vary?

A. Depends how Conflict is defined.


L is Lipschitz constant: $f(x)-f(y) < L |x-y|$

Four common
methods
in Graphics


Method	Distance Function	Order Independent	Full Coverage	Conflict Free	Edge Min	Edge Max	Sin Angle Min	Max L
Prior	$r(\mathbf{x})$	no	no	no	$1/(1+L)$	$2/(1-2L)$	$(1-2L)/2$	$1/2$
Current	$r(\mathbf{y})$	no	no	no	$1/(1+L)$	$2/(1-L)$	$(1-L)/2$	1
Bigger	$\max(r(\mathbf{x}), r(\mathbf{y}))$	yes	no	yes	1	$2/(1-2L)$	$(1-2L)/2$	$1/2$
Smaller	$\min(r(\mathbf{x}), r(\mathbf{y}))$	yes	yes	no	$1/(1+L)$	$2/(1-L)$	$(1-L)/2$	1




Prior:
new candidate disk center inside an old prior disk




Current:
old prior disk center inside a new candidate disk



Bigger:
small disk center inside big disk center






Smaller:
big disk center inside small disk center



Bigger is stricter than
Sphere packing:
 $\frac{1}{2}$ radius disks overlap
distance: $\text{sum}(r(x), r(y))/2$



Graphics Quality Criteria



Point Set Analysis: <http://code.google.com/p/psa/>

OK, but what about these?

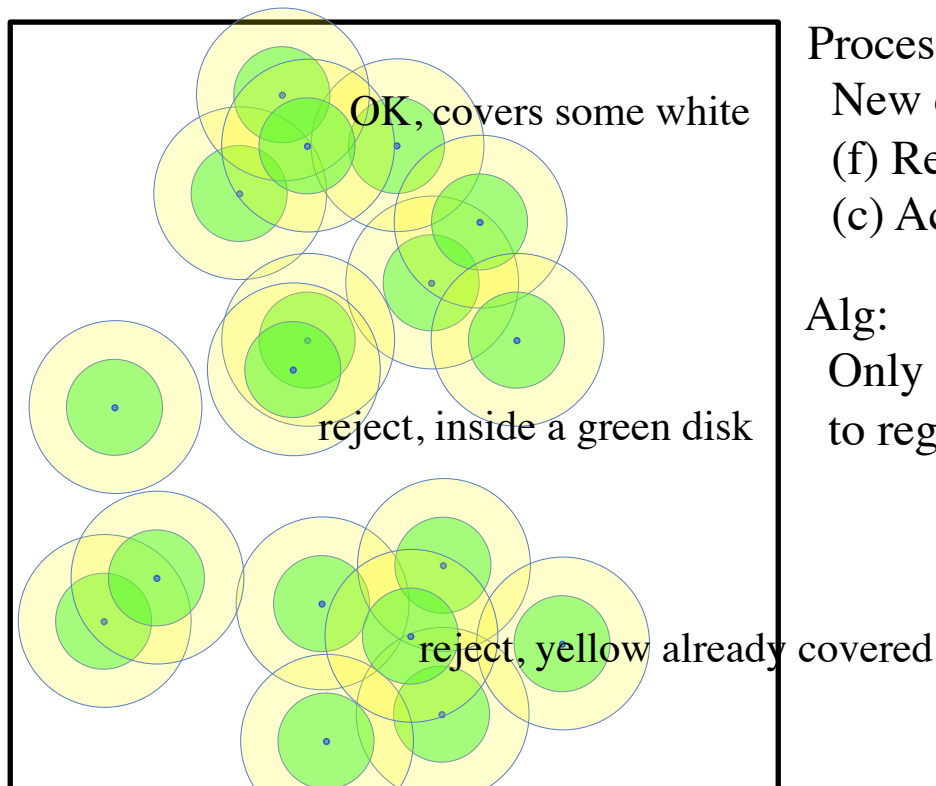
- **Graphics papers say they like MPS because**
 - **Small low frequency component**
 - **No big spikes, especially spikes at high frequency**
 - **IMO want truncated white-noise**

Unknown: analytic description of the limit distribution for MPS,
Mean location and magnitude, std deviations of peaks?
Anyone know some good spatial statisticians to work with?



Our Solution (second contribution)

- Disk coverage radius larger than free radius
 $R_c > R_f$ (yellow > green)
- New disks must cover some unique uncovered area
 - Else maximal (limit) distribution would be the same
 - Contrast to Hard-core Strauss disc process:
coverage disks are observed, no effect on process



Process:

New candidate point uniform at random

(f) Rejected if center inside a small green disk

(c) Accepted if its yellow disk covers some white area

Alg:




Only generate points in an outer approximation to regions satisfying (c) and (f) in the first place.



Two-radii MPS output

- **Classic MPS**




$$R_f = R_c$$



- **Two-radii MPS**

$$2 R_f = R_c$$




- $R_f = \text{min center dist}$
- $R_c = \text{max Vor dist}$



- **Uniform**

$$R = 0$$



non-maximal





Random refinements by shrinking radius

Continuously shrink radii



Coverage radius (blue) larger than inhibition radius (red)


With one radius, get deterministic point placement,
at Voronoi vertex, as classic Delaunay Refinement

With two radii, random placement,
neighborhood of Voronoi vertex: inside green disk-at-u outside red




Hierarchical by shrinking radius


Discrete steps



(a) $t = 0.8$ end



(b) $t = 0.6$ start



(c) $t = 0.6$ end


Simple since continue original algorithm

No noticeable effect on output spectrum



Contrast to Delaunay Refinement (DR)







- ***Direct sizing control***
 - ***Disks centered at points, sized at that point***
- ***Never $O(n^2)$ – no intermediate triangulation, generate points first***
- ***No concentration at medial axis as DR***
- ***Global Random placement***
 - ***Slow vs. local deterministic DR***
 - ***Spectrum results for DR depends on target, queue order***
Alex Rand experiments in progress







Spectrum results for DR

- Depends on target, queue order
(Alex Rand experiments in progress)





Uniform MPS vs. DR angles and edges



To do: study and contrast further



Thanks!

